Optimal, Robust Predictive Control of Nonlinear Systems under Probabilistic Uncertainty using Particles

Lars Blackmore and Brian C. Williams

***Abstract*— In this paper we present a novel method for robust, optimal control of nonlinear systems under probabilistic uncertainty. The method extends a previous approach for linear systems that approximates the distribution of the predicted system state using a finite number of particles. We couple this particle-based approach with a nonlinear solver that does not take into account uncertainty to give a new method for nonlinear, robust control. Any solution returned by the algorithm is guaranteed to be ϵ-close to a local optimum of the nonlinear stochastic control problem.**

I. Introduction

Autonomous vehicles need to be able to plan trajectories to a specified goal that avoid obstacles, and are robust to the inherent uncertainty in the problem[3], [14], [15], [16], [17]. This uncertainty arises due to uncertain state estimation, disturbances and modeling errors, which can be described using probabilistic models of uncertainty.

Previous work posed this problem as *chance-constrained* optimal predictive control. Chance constraints ensure the probability of failure is less than a specified threshold. A number of solutions exist for the case where all uncertain variables have Gaussian distributions, and the system dynamics are linear[1], [3], [7], [12], [18]. In the case of non-Gaussian distributions, our previous work solved the problem by approximating the distribution of the system using a finite number of particles[2]. In this manner an intractable stochastic control problem can be approximated as a deterministic one, and for linear systems the resulting optimization can be solved to global optimality extremely efficiently using Mixed-Integer Linear Programming[8]. Furthermore, as the number of particles tends to infinity, the approximation becomes exact.

While in estimation particle-based approaches have been shown to give the most benefit with nonlinear systems and non-Gaussian noise, the particle control approach was, however, limited to linear systems. In this paper, we extend the approach to apply to nonlinear systems. The key idea behind the approach is to use an initial solution to the nonlinear control problem that does not take into account uncertainty, and then use the particle control approach to turn this non-robust solution into a robust one. Any solution returned by the algorithm is guaranteed to be ϵ-close to a local optimum of the nonlinear stochastic control problem.

L. Blackmore is a PhD candidate at the Massachusetts Institute of Technology, Cambridge, MA 02139. larsb@mit.edu

B. C. Williams is an associate professor with the department of Aeronautics and Astronautics at the Massachusetts Institute of Technology, Cambridge, MA 02139. williams@mit.edu

II. Problem Statement

We consider a discrete-time dynamic system where the future states $\mathbf{x}_{1:T} \triangleq [\mathbf{x}_1, \cdots, \mathbf{x}_T]$ are functions of the control inputs $\mathbf{u}_{0:T-1} \triangleq [\mathbf{u}_0, \cdots, \mathbf{u}_{T-1}]$, the initial state $\mathbf{x}_0$, and disturbances $\nu_{0:T-1} \triangleq \nu_0, \cdots, \nu_{T-1}$:

$$\begin{aligned} \mathbf{x}_1 &= f_1(\mathbf{x}_0, \mathbf{u}_0, \nu_0) \\ \mathbf{x}_2 &= f_2(\mathbf{x}_0, \mathbf{u}_{0:1}, \nu_{0:1}) \\ &\vdots \\ \mathbf{x}_T &= f_T(\mathbf{x}_0, \mathbf{u}_{0:T-1}, \nu_{0:T-1}). \end{aligned} \tag{1}$$

The initial state and disturbances are uncertain, but are modeled as random variables with known distributions. Hence the future states are also random variables, whose distributions depend on the control inputs. We assume that the initial state and disturbances are independent. Modeling errors can be modeled as an additional stochastic disturbance process[13][1]. Note that we make *no assumptions* on the form of the uncertain distributions.

In this paper we are concerned with the following stochastic predictive control problem:

> *Choose a finite, optimal sequence of control inputs from a continuous, bounded set U, which takes into account probabilistic uncertainty so that the probability of the system state leaving a given feasible region F is at most δ.*

Optimality is defined in terms of minimizing $h(\mathbf{u}_{0:T-1}, \mathbf{x}_{1:T})$. This function can be defined in terms of minimizing control effort, for example. In the case of vehicle path planning, the feasible region can be defined so that the system state is in a goal region at the final time step, and avoids a number of obstacles at all time steps.

III. Review of Particle Control for Linear Systems

In [2] we described an approach for solving the problem in Section II. The key idea is that by approximating all probabilistic distributions using particles, an intractable stochastic optimization problem can be approximated as a tractable deterministic one. The particle-based approach enables us to deal with arbitrary probability distributions. By solving this deterministic problem we obtain an approximate solution to the original stochastic problem, with the additional property

[1]For notational simplicity we assume for the rest of the development a single disturbance process; however the method applies equally to multiple disturbance processes, and hence to modeling errors as well as external disturbances.

that as the number of particles used tends to infinity, the approximation becomes exact. The approach is summarized as follows:

1) Generate N samples from the joint distribution of initial state and disturbances.
2) Express the distribution of the future state trajectories approximately as a set of N *analytic* particles, where each particle $\mathbf{x}^{(i)}_{1:T}$ corresponds to the state trajectory given a particular set of samples. Each particle *depends explicitly on the control inputs* $\mathbf{u}_{0:T-1}$, which are yet to be generated.
3) Approximate the chance constraints in terms of the generated particles; the probability of $\mathbf{x}_{1:T}$ falling outside of the feasible region is approximated as the *fraction* of particles $\mathbf{x}^{(i)}_{1:T}$ that do so.
4) Approximate the cost function in terms of particles.
5) Solve the deterministic constrained optimization problem for control inputs $\mathbf{u}_{0:T-1}$.

The general particle control problem results in a deterministic optimization problem that is intractable, except for very small problems. However in [2] we showed that for a polytopic feasible region F, piecewise linear cost function h and linear system dynamics $\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t$, the deterministic optimization can be solved to global optimality in an efficient manner using Mixed-Integer Linear Programming. This relies on the fact that each particle $\mathbf{x}^{(i)}_{1:T}$ is a linear function of the control input sequence $\mathbf{u}_{1:T-1}$.This is also true for time-varying linear systems. We now couple the particle control approach for time-varying linear systems with a non-robust planner for nonlinear systems to give a tractable particle control approach for nonlinear systems.

IV. Particle Control for Nonlinear Systems

The key idea behind the new approach is to use the particle control method described in Section III to create a solution to the robust nonlinear control problem, starting from a solution generated by a non-robust nonlinear controller. The robust solution search is limited to a linearization region around the non-robust solution. If this does not contain a robust solution, a new non-robust plan is generated with the linearization region removed from the search space. By repeating this process, the search space is explored to find robust solutions even in the case of non-convex problems. We assume that we have a nonlinear planner that, given a given a nonlinear, discrete-time, deterministic planning problem, and a feasible region that *may be nonconvex*, returns a locally optimal control sequence, which we denote $\bar{\mathbf{u}}_{1:T}$. Examples of such planners include[5], [6], [9], [10], [11] The new algorithm is given in Table I and illustrated in Figure 1.

Lemma 1. *Any solution returned by the new method is ϵ-close to a local optimum of the nonlinear stochastic control problem.*

Proof: The new approach returns a solution $\mathbf{u}^*_{1:T}$ only if a robust solution to the linearized particle control problem was found in the interior of linearization region U^k. This plan is the optimal linearized particle control solution in region U^k, which would be a local optimum to the nonlinear robust problem except for the approximation error introduced by linearization and sampling. Linearization approximates the system dynamics and cost function, but since the linearization region is known and bounded, the approximation error can be bounded using standard results. Particles approximate the true distribution of the uncertain variables, but results from prior work in particle filtering exist that bound the approximation here also[4]. Hence by specifying the size of the linearization region and the number of particles used, the error between the solution $\mathbf{u}^*_{1:T}$ and the locally optimal solution to the nonlinear problem can be bounded to a desired value ϵ. □

1) *Initialization.* Assign $k = 1$, $U^1 = U$ and $h^* = +\infty$
2) *Generate the non-robust control problem* $\bar{P}^k = \langle \bar{f}, F, U^k, \bar{h}\rangle$. The deterministic system dynamics $\bar{f}$ and cost function $\bar{h}$ are generated by setting the uncertain variables $\mathbf{x}_0$ and ν_t to their mean values. The deterministic system state is constrained to stay within F.
3) *Solve the non-robust control problem* $\bar{P}^k$. This gives the globally optimal control sequence $\bar{\mathbf{u}}^k_{0:T-1}$, the corresponding state sequence $\bar{\mathbf{x}}^k_{1:T}$, and the cost $\bar{h}^k$ for iteration k. If the problem is infeasible, stop.
4) *Check for existing robust solution*. If $h^* \leq \bar{h}^k$, stop.
5) *Generate fixed particles*. Generate N samples of the initial state distribution and disturbances. Calculate the state trajectory for each particle assuming that the control input is *fixed* at $\bar{\mathbf{u}}^k_{0:T-1}$. We denote these fixed particle trajectories as $\hat{\mathbf{x}}^{(i)}_{1:T}$:
$$\hat{\mathbf{x}}^{(i)}_{1:T} = \left[\hat{\mathbf{x}}^{(i)\prime}_1\ \hat{\mathbf{x}}^{(i)\prime}_2\ \cdots\ \hat{\mathbf{x}}^{(i)\prime}_T\right]' \qquad (2)$$
$$\hat{\mathbf{x}}^{(i)}_t = f_t(\mathbf{x}^{(i)}_0, \bar{\mathbf{u}}^k_{0:t-1}, \nu^{(i)}_{0:t-1}).$$
6) *Linearize for each particle.* For each fixed particle, linearize the system dynamics about the state trajectory $\hat{\mathbf{x}}^{(i)}_{1:T}$ to generate a time-varying linear system $\langle A^{(i)}_t, B^{(i)}_t\rangle$.
7) *Calculate linearized analytic particles.* For each particle, using the same samples for the uncertain variables as generated in Step 5, calculate the state trajectory *as a function of the control inputs* $\mathbf{u}_{0:T-1}$:
$$\mathbf{x}^{(i)}_{1:T} = \left[\mathbf{x}^{(i)\prime}_1\ \mathbf{x}^{(i)\prime}_2\ \cdots\ \mathbf{x}^{(i)\prime}_T\right]'$$
$$\mathbf{x}^{(i)}_t = \prod_{l=1}^{t} A^{(i)}_l \mathbf{x}^{(i)}_0 + \sum_{j=0}^{t-1}\Big(\prod_{l=1}^{t-j-1} A^{(i)}_l\Big) B^{(i)}_j(\mathbf{u}_j + \nu^{(i)}_j). \quad (3)$$
8) *Solve linearized particle control problem* P^k. Each particle $\mathbf{x}^{(i)}_{1:T}$ is a linear function of the control inputs $\mathbf{u}_{0:T-1}$. The cost function h is linearized about the non-robust control solution $\bar{\mathbf{u}}^k_{0:T-1}$. Using these linearized analytic functions, the problem of finding control inputs such that $p(\mathbf{x}_{1:T} \notin F) \leq \delta$ is approximated using the particle control approach described in Section III. We add the constraint that the control input must lie close to $\bar{\mathbf{u}}_{0:T-1}$, within a polytopic linearization region R, the size of which is specified by the user. The problem is solved to give a robust solution $\mathbf{u}^k_{0:T-1}$ with cost h^k, or is infeasible.
9) *Remove linearization region.* Increment k and assign $U^k = U^{k-1}/R$ where the operator $/$ denotes the set difference. If a feasible solution $\mathbf{u}^k_{0:T-1}$ was found in the strict interior of U^k, assign $h^* = \min\{h^k, h^*\}$.
10) Go to step 2.

TABLE I. Particle Control Approach for Nonlinear Systems

V. CONCLUSION

We have presented a novel approach to robust, optimal control of nonlinear systems under uncertainty. Chance constraints are used to formulate robustness. The new method extends our previous particle-based approach for linear systems by linearizing around a solution generated by a nonlinear solver that does not take into account uncertainty. Any solution returned by the new method is guaranteed to be ϵ-close to a local optimum of the nonlinear stochastic control problem. Future work will validate this approach empirically and will develop extensions of the method that guarantee global optimality and completeness.

REFERENCES

[1] I. Batina. *Model Predictive Control for Stochastic Systems by Randomized Algorithms*. PhD thesis, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2004.

[2] L. Blackmore. A probabilistic particle control approach to optimal, robust predictive control. In *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2006.

[3] L. Blackmore, H. X. Li, and B. C. Williams. A probabilistic approach to optimal robust path planning with obstacles. In *Proceedings of the American Control Conference*, 2006.

[4] D. Fox. Adapting the sample size in particle filters through kld-sampling. *International Journal of Robotics Research (IJRR)*, 22, 2003.

[5] E. Frazzoli, M. Dahleh, and E. Feron. Real-time motion planning for agile autonomous vehicles. In *Proceedings of AIAA Conference on Guidance, Navigation and Control*, 2000.

[6] M. Henson. Nonlinear model predictive control: current status and future directions. *Computers and Chemical Engineering*, 23, 1998.

[7] D. H. Van Hessem. *Stochastic Inequality Constrained Closed-loop Model Predictive Control*. PhD thesis, Technische Universiteit Delft, Delft, The Netherlands, 2004.

[8] ILOG. Ilog cplex user's guide, 1999.

[9] L. E. Kavraki, P. Svestka, J. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4), 1996.

[10] D. Kogan. Realtime path planning via nonlinear optimization methods. Master's thesis, California Institute of Technology, Pasadena, CA, 2005.

[11] J. Kuffner, K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue. Motion planning for humanoid robots. In *Proceedings of 20th International Symposium on Robots Research*, 2003.

[12] P. Li, M. Wendt, and G. Wozny. A probabilistically constrained model predictive controller. *Automatica*, 38:1171–1176, 2002.

[13] L. Ljung. *System Identification: Theory for the user*. Prentice Hall, New York, 1999.

[14] A. Y. Ng and M. Jordan. Pegasus: A policy search method for large mdps and pomdps. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, 2000.

[15] J. Pineau and G. Gordon. Pomdp planning for robust robot control. In *12th International Symposium of Robotics Research*, 2005.

[16] J. Pineau, G. Gordon, and S. Thrun. Policy-contingent abstraction for robust robot control. In *Proceedings of the Conference on Uncertainty in AI (UAI)*, Acapulco, Mexico, 2003.

[17] N. Roy, G. Gordon, and S. Thrun. Planning under uncertainty for reliable health care robotics. In *Proceedings of the International Conference on Field and Service Robotics*, 2003.

[18] A. Schwarm and M. Nikolaou. Chance-constrained model predictive control. *AlChE Journal*, 45:1743–1752, 1999.

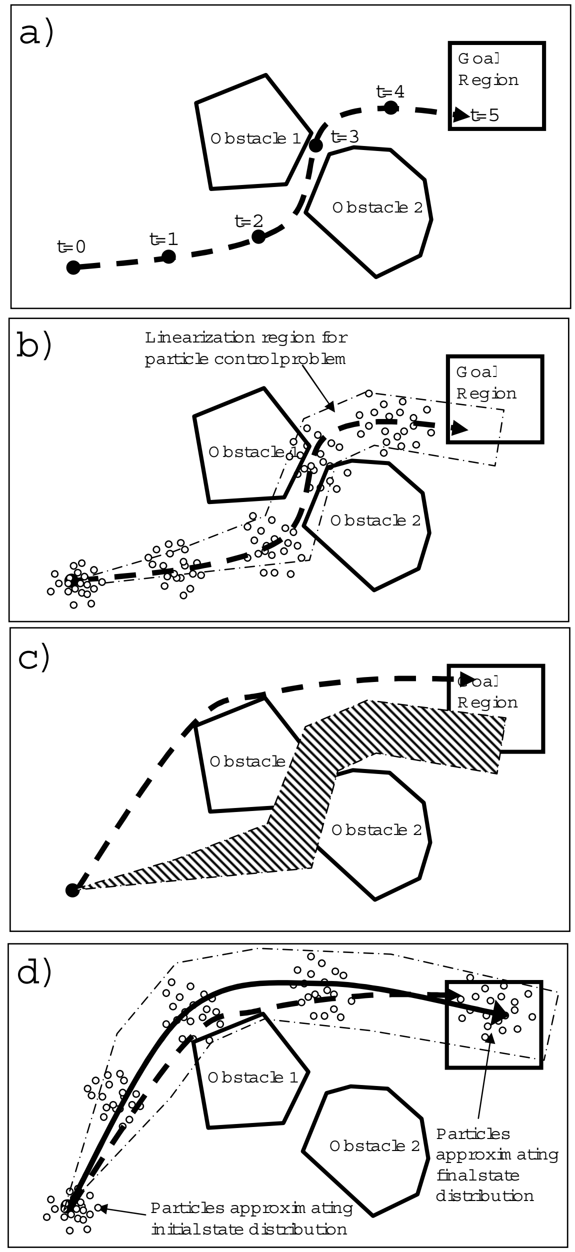

Fig. 1. Illustration of new nonlinear particle control approach for obstacle avoidance example. The system state must avoid obstacles and end in the goal region with probability at least 90%. a) The non-robust problem is solved using a nonlinear discrete-time planner that does not take into account uncertainty. b) A particle control problem is generated by linearizing about the non-robust solution. In this case no robust solution exists within the linearization region. c) The nonlinear planner finds a new solution that also avoids the linearization region. d) A new particle control problem is generated about the new non-robust solution. This problem is solved to find a robust solution. In this case a robust solution exists, and exactly 90% of the particles succeed.